Searching for periodic sources with LIGO

Patrick R Brady[(1)], Teviet Creighton[(1)], Curt Cutler[(2),(3)] and Bernard F Schutz[(3)]
[(1)] *Theoretical Astrophysics 130-33, California Institute of Technology, Pasadena, CA 91125*
[(2)] *Center for Gravitational Physics and Geometry, Pennsylvania State University, University Park, PA 16802*
[(3)] *Max Planck Institute for Gravitational Physics, Albert Einstein Institute, Schlaatzweg 1, D-14473 Potsdam*
(February 7, 2008)

arXiv:gr-qc/9702050v1 25 Feb 1997

We investigate the computational requirements for all-sky, all-frequency searches for gravitational waves from spinning neutron stars, using archived data from interferometric gravitational wave detectors such as LIGO. These sources are expected to be weak, so the optimal strategy involves coherent accumulaton of signal-to-noise using Fourier transforms of long stretches of data (months to years). Earth-motion-induced Doppler shifts, and intrinsic pulsar spindown, will reduce the narrow-band signal-to-noise by spreading power across many frequency bins; therefore, it is necessary to correct for these effects before performing the Fourier transform. The corrections can be implemented by a parametrized model, in which one does a search over a discrete set of parameter values (*points* in the parameter space of corrections). We define a metric on this parameter space, which can be used to determine the optimal spacing between points in a search; the metric is used to compute the number of independent parameter-space points N_p that must be searched, as a function of observation time T. This method accounts automatically for correlations between the spindown and Doppler corrections. The number $N_p(T)$ depends on the maximum gravitational wave frequency and the minimum spindown age $\tau = f/\dot{f}$ that the search can detect. The signal-to-noise ratio required, in order to have 99% confidence of a detection, also depends on $N_p(T)$. We find that for an all-sky, all-frequency search lasting $T = 10^7$ s, this detection threshhold is $h_c \approx (4 - -5)h_{3/\mathrm{yr}}$, where $h_{3/\mathrm{yr}}$ is the corresponding 99% confidence threshhold if one knows in advance the pulsar position and spin period.

We define a coherent search, over some data stream of length T, to be one where we apply a correction, followed by an FFT of the data, for every independent point in the parameter space. Given realistic limits on computing power, and assuming that data analysis proceeds at the same rate as data acquisition (e.g. 10 days of data gets analysed in ~ 10 days), we can place limitations on how much data can be searched coherently. In an all-sky search for pulsars having gravity-wave frequencies $f \leq 200$Hz and spindown ages $\tau \geq 1000$Yrs, one can coherently search ~ 18 days of data on a teraflops computer. In contrast, a teraflops computer can only perform a ~ 0.8-day coherent search for pulsars with frequencies $f \leq 1$kHz and spindown ages as low as 40Yrs.

In addition to all-sky searches we consider coherent directed searches, where one knows in advance the source position but not the period. (Nearby supernova remnants and the galactic center are obvious places to look.) We show that for such a search, one gains a factor of ~ 10 in observation time over the case of an all-sky search, given a 1Tflops computer.

The enormous computational burden involved in coherent searches indicates a need for alternative data analysis strategies. As an example we briefly discuss the implementation of a simple hierarchical search in the last section of the paper. Further work is required to determine the optimal approach.

I. INTRODUCTION

The direct observation of gravitational waves is a realistic goal for the kilometer-scale interferometers which are now under construction at various sites around the world [1,2]. However, the battle to see these waves is not over when the detectors are constructed and running. Searching for gravitational wave signals in the interferometer output presents its own problems, not least of which is the sheer volume of data involved.

Potential sources of gravitational waves fall roughly into three classes: bursts, stochastic background, and continuous emitters.

Burst sources produce signals which last for times considerably shorter than available observation times. The chirp signals from compact coalescing binaries belong to this class. Since theoretical waveforms, valid during the inspiral phase of the binary evolution, have been accurately calculated using post-Newtonian methods [3], it is possible to search the data stream for chirps using matched filtering techniques. Detailed studies have been carried out to ascertain the optimal set of search templates [4,5], and a preliminary investigation of search algorithms is now under way [6]. Detection of other, not so well understood, sources in this class—e.g. non-axisymmetric supernovae—has received limited attention [7].

Flanagan [8] has determined how to cross-correllate the output of two detectors in order to search for a stochastic background of gravitational radiation, which was implemented by Compton [9] and applied to data taken during a period of 100 hours by two prototype interferometers detectors in Glasgow and Garching [10]. In [11], Allen presents a detailed discussion of the potential significance of detecting a stochastic background. Compton's work, and simulations performed by Allen, have demonstrated that this kind of analysis requires minimal computational resources.

In this paper we consider some issues involved in searching for continuous wave sources. Throughout our discussion we use pulsars as a guide to develop a search strategy.

A. Gravitational waves from pulsars

Rapidly rotating neutron stars (pulsars) tend to be axisymmetric; however, they must break this symmetry in order to radiate gravitationally. The pulsar literature contains several mechanisms which may lead to deformations of the star, or to precession of its rotation axis, and hence to gravitational wave emission. The characteristic amplitude[1] of gravitational waves from pulsars scales as

$$h_c \sim \frac{If^2\epsilon}{r} \tag{1.1}$$

where I is the moment of inertia of the pulsar, f is the gravitational wave frequency, ϵ is a measure of the deviation from axisymmetry and r is the distance to the pulsar.

Pulsars are thought to form in supernova explosions. The outer layers of the star crystallize as the newborn pulsar cools by neutrino emission. Estimates, based on the expected breaking strain of the crystal lattice, suggest that anisotropic stresses, which build up as the pulsar loses rotational energy, could lead to $\epsilon \lesssim 10^{-5}$; the exact value depends on the breaking strain of the neutron star crust as well as the neutron star's 'geological history', and could be several orders of magnitude smaller. Nonetheless, this upper limit makes pulsars a potentially interesting source for kilometer scale interferometers. Figure 1 shows some upper bounds on the amplitude due to these effects.

Large magnetic fields trapped inside the superfluid interior of a pulsar may also induce deformations of the star. This mechanism has been explored recently in [12], indicating that the effect is extremely small for standard neutron star models ($\epsilon \lesssim 10^{-9}$).

Another plausible mechanism for the emission of gravitational radiation in very rapidly spinning stars is the Chandrasekhar-Friedman-Schutz (CFS) instability, which is driven by gravitational radiation reaction [13,14]. It is possible that newly-formed neutron stars may go through this instability spontaneously as they cool soon after formation. The radiation is emitted at a frequency determined by the frequency of the unstable normal mode, which may be less than the spin frequency.

Accretion is another way to excite neutron stars into emitting gravitational waves. Wagoner [15] proposed that accretion may drive the CFS instability. There is also the Zimmermann-Szedinits mechanism [16] where the principal axes of the moment of inertia are driven away from the rotational axes by accretion from a companion star. Accretion can in principle produce relatively strong radiation, since the amplitude is related to the accretion rate rather than to structural effects in the star. However, accreting neutron stars will be in binary systems, and these present problems for detection that go beyond the ones we discuss in this paper. We hope to return to the problem of looking for radiation from orbiting neutron stars in a future publication.

[1] We adopt the definition of h_c provided in Eq. (50) of Thorne [7].

B. Three classes of sources

Observed pulsars fall roughly into two groups: (i) young, isolated pulsars having periods of tens or hundreds of milliseconds, and (ii) older, millisecond pulsars. The young pulsars are most likely to deviate significantly from axisymmetry; however, they are generally observed to have low frequencies, so that there is a competition between the frequency, f, and deviation from axisymmetry, ϵ, in Eq. (1.1). On the other hand, millisecond pulsars, whose waves are higher in frequency, tend to be quite old and well annealed into an axisymmetric configuration.

Radio observations can only probe a small portion of our galaxy in searching for pulsars. A significant effect reducing the depth of radio searches is dispersion of the signal by galactic matter between potential sources and the earth. Given current evolutionary scenarios for pulsars—that they are born in supernova explosions—it seems likely that most pulsars should be located in the galactic disk, and the youngest of these will also be shrouded in a supernova remnant, making them invisible to radio astronomers.

Blandford [17,7] has pointed out that there could exist a class of pulsars which spin down primarily due to gravitational radiation reaction. For sources in this class the frequency scales as $f \propto \tau^{-1/4}$, where τ is the age of the pulsar. If the mean birth rate for such pulsars in our galaxy is τ_B^{-1}, the nearest one should be a distance $r = R\sqrt{\tau_B/\tau}$ from earth, where $R \simeq 10$kpc is the radius of the galaxy. The intrinsic gravitational wave amplitude (that is, the amplitude h at some fixed distance) of a pulsar in this class is proportional to $\tau^{-1/2}$. Thus, the nearest source in this class would have a dimensionless amplitude h_c at the Earth

$$h \simeq 8 \times 10^{-25} \left(\frac{200\text{Yrs}}{\tau_B}\right)^{1/2} \tag{1.2}$$

independent of the frequency, or the the ellipticity ϵ of the pulsar. Assuming the existence of such a class of pulsars, with $\tau_B \lesssim 2 \times 10^4$ Yrs, we see from Fig. 1 that there is a large region of parameter space that is both (i) detectable by the LIGO detector and (ii) physically reasonable, in the sense that $\epsilon < 10^{-5}$ and f lies in the range 200–1000 Hz.

Note that Blandford's argument can be slightly re-cast to yield an upper limit on the gravitational wave strength of *any* isolated pulsar–i.e., any pulsar whose radiated angular momentum is not being replenished by accretion. The age of an isolated pulsar must be shorter than the age computed assuming the spin-down is solely due to gravitational wave emission. Correspondingly, if we set τ_B equal to 40 yrs (corresponding to the birthrate for *all* pulsars), we get the following upper limit for measured gravitational wave amplitude of an isolated pulsar: $h_c < 2 \times 10^{-24}$. Of course, this is a statistical argument. This bound could certainly be violated by an isolated pulsar that just happens to be anomalously close to us.

It is important that any search strategy should be general enough to encompass all three of the above classes, allowing for the significant changes in frequency which may be inherent in the sources (see section II).

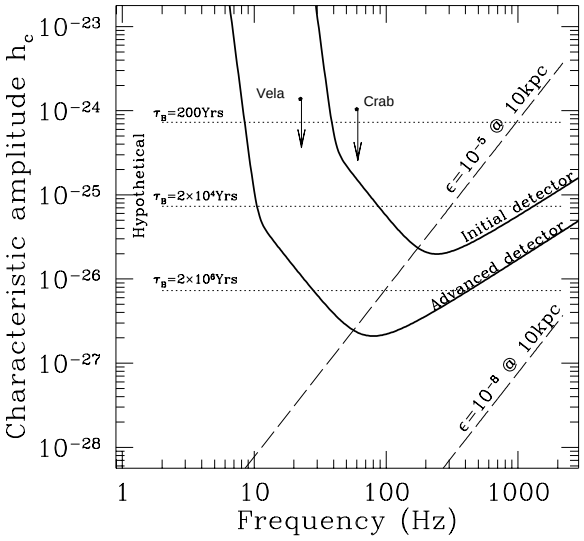

FIG. 1. Characteristic amplitudes h_c [see Eq. (3.5)] for several postulated periodic sources, compared with sensitivities $h_{3/\text{yr}}$ of the initial and advanced detectors in LIGO. ($h_{3/\text{yr}}$ corresponds to the amplitude h_c of the weakest source detectable with 99% confidence in $\frac{1}{3}$yr $= 10^7$s integration time, if the frequency and phase of the signal is known in advance.) Long-dashed lines show the expected signal strength as a function of frequency for pulsars at a distance of 10 kpc, assuming non-axisymmetries of $\epsilon = 10^{-5}$ and $\epsilon = 10^{-8}$, where ϵ is defined in section IIIA. Upper limits are also plotted for the Crab and Vela pulsars, assuming their entire measured spin-down is due to gravitational wave emission. The dotted lines indicate the strongest waves received at the earth for Blandfords hypothetical class of pulsars; each line corresponds to a particular birth rate.

C. The data analysis problem

The detection of continuous, nearly fixed frequency waves will be achieved by constructing power spectrum estimators and searching for statistically significant peaks at fixed frequencies. In practice, this is achieved by calculating the amplitude of the Fourier transform of the detector output given by applying a fast Fourier transform (FFT), a discrete approximation to the true Fourier transform:

$$\tilde{h}(f) = \frac{1}{\sqrt{T}} \int_0^T e^{2\pi i f t}\, h(t)\, dt\,. \tag{1.3}$$

The main hope of detection lies in the fact that one may observe the sky for long time periods of time T. When such a data stretch is transformed to make the underlying

signal monochromatic, the signal to noise ratio grows as $\sqrt{T}$ in amplitude (or as T in the power spectrum). One will likely need to have integration times of several weeks or months in order for the expected signals from nearby sources to rise above the noise. However, such long data stretches pose a significant computational burden; using 10^7s of data to look for signals with gravitational wave frequencies up to 500Hz requires calculating an FFT with $N \simeq 10^{10}$ data points. Calculation of a single such FFT would take about 1s on a 1Tflops computer, assuming that all 10^{10} points can be held in fast memory. Unfortunately, this is not the whole story.

The detection problem is complicated by the fact that the signal received at the detector is not perfectly monochromatic. Earth-bound detectors participate in complex motions which lead to significant Doppler shifts in frequency as the Earth rotates, and as it orbits around the sun (this orbit is significantly perturbed by the moon and the other planets). The time dependent accelerations broaden the spectral lines of fixed frequency sources spreading power into many Fourier bins about the observed frequency. In order to maintain the benefit of long observation times, it is therefore necessary to remove the effects of the detector motion from the data stream. This can achieved by introducing an inertial (barycentered) time coordinate and carrying out the FFT with respect to it. The difficulty of doing this was first estimated by one of us [18]. However, we must also consider the additional complication that the signal may not be intrinsically monochromatic. If the signal exhibits intrinsic frequency drift, or modulation, due to the nature and location of the source — as is expected for pulsars which spin down with time — these effects can also be removed in the transformation to the new time coordinate.

Unfortunately, the demodulated time coordinate depends strongly on the direction from which the signal is expected, and on the intrinsic frequency evolution one assumes for the source. Thus, in searching for sources whose position and timing are not well known in advance one must apply many different corrections to the data, performing a new FFT after each correction. Given the possibility that the strongest sources of continuous gravitational waves may be electromagnetically invisible or previously undiscovered, an *all sky, all frequency* search for such unknown sources is of considerable interest. To obtain some idea of the magnitude of this task, consider searching the entire sky for signals with (fixed) frequencies up to 500Hz using 10^7s worth of data. Assuming the entire data stream could be held in fast memory on a machine capable of 1Tflops, it would take 10^8s to complete the search. Introducing intrinsic spindown effects into the search increases the computational cost, at fixed integration time, by many orders of magnitude. This computational cost is the central problem of searching for unknown pulsars in the output from gravitational wave detectors and is the focus of this paper.

D. Summary of results

We parametrise the space of pulsar signals by the position of the source on the sky $\{\theta, \phi\}$, entering through Doppler shifts due to the detector's motion, and by spindown parameters f_k which characterise the intrinsic frequency evolution. [See Eq. (3.7).] We constrain the range of possible values of the spindown parameters using the (spindown) age $\tau = f/\dot{f}$ of the youngest search that a search can detect, thus $|f_k| \leq \tau^{-k}$. For the computationally-intensive search over all sky positions and spindown parameters, it is important to be able to calculate the smallest number of independent parameter values which must be sampled in order to cover the entire space of signals. We have accomplished this by introducing a distance measure and corresponding metric on the parameter space. The analysis is patterned after a similar one developed by Owen [5] for gravitational waves from inspiralling, compact binaries. Using our metric one can compute the volume of parameter space, thus inferring the number of independent points that must be sampled in order to cover the entire space. We define a *coherent search* to be one where we perform one demodulation and FFT of the data for every independent point in the parameter space. Besides telling us the computational requirements for a coherent search, the metric approach tells us how to place the points most efficiently in parameter space, in a similar way to that discussed by Owen.

We have found it useful to present the results based on several possible search strategies, which cover different regions of the parameter space. Accordingly, we define a pulsar to be *old* if its spindown age τ is greater than 10^3 Yrs and *young* if $\tau \gtrsim 40$ Yrs. A pulsar is considered to be *slow* if its gravitational wave frequency is $f \lesssim 200$ Hz and *fast* if $f \lesssim 10^3$ Hz.

A coherent all-sky search of 10^7 seconds of data for old, slow pulsars requires approximately 1.1×10^{10} independent points in the parameter space; only one spindown parameter is needed to account for intrinsic frequency evolution. In contrast, an all-sky search for fast, young pulsars in 10^7 seconds of data requires 8×10^{21} independent parameter space points to be sampled, using three spindown parameters to model intrinsic frequency evolution. Note that searches for old, fast pulsars (such as known millisecond radio pulsars) and young, slow pulsars (younger brothers of the Crab and Vela) are automatically subsumed under the latter search. These results mean the following. Assuming unlimited computer power and stationary, gaussian statistics, a pulsar with unknown position and period must have strain $h_c \approx 4.3 h_{3/\mathrm{yr}}$, if it is in our 'old, slow' category, and $h_c \approx 5.1 h_{3/\mathrm{yr}}$, if it is in our 'young, fast' category, to be detected with 99% confidence in a 10^7 second search. Here $h_{3/\mathrm{yr}}$ is the strain required for detection with 99%-confidence in a 10^7 second integration, assuming the pul-

sar position and period are known in advance[2]:

$$h_{3/\mathrm{yr}}(f) = 4.2\,\sqrt{S_n(f) \times 10^{-7} Hz}\,. \qquad (1.4)$$

Thus, when considering an all-sky, all-frequency pulsar search, the LIGO sensitivity curves shown in Fig.1 effectively overestimate the detector's sensitivity by a factor of $\sim 4--5$, even in the limit of infinite computing power.

Our ability to perform searches for continuous waves will certainly be limited by the available computing resources. Assuming realistic computer power — say of order 10^{13} flops — we estimate that computing limitations will effectively reduce the sensitivity of the detector by another factor of ~ 2, even for some reasonably optimized and efficient search strategy. However more work will be needed to develop a reasonably optimized algorithm, and thus to refine this latter estimate.

While the concept of the metric is introduced in the framework of an all-sky search for unknown pulsars, it is clear that we may use the same approach to examine the depth of a search over limited regions of the parameter space. In particular, once the scope of a search is decided, the optimization procedure dicussed in section VI can be used to determine the observation time and grid spacing which maximises the expected sensitivity of a search. As an example, we consider coherent *directed searches*, in which one assumes a specific sky position (such as a particular cluster or supernova remnant) and searches only over spindown parameters. Again, we present results for two concrete scenarios based on fast, young pulsars and old, slow pulsars. Similar considerations apply to directed searches as to all-sky searches; that is, the curves in Fig. 1 overestimate the detector sensitivity for 10^7 second integration. Table 1 summarises the results for both cases.

We note that in each type of search, the number of parameter space points, and hence the computational requirements, were reduced significantly by the assumption that the points were placed with optimal spacings given by the metric formalism. Nevertheless, the bottom line is that limitations on computational resources will severely restrict the integration times that can be achieved. Assuming access to a Tflops of computing power (effective computational throughput, ignoring possible overheads due to interprocessor communication or data access), we find the following limits on coherent integration times: For young, fast pulsars we are limited to about 0.8 days for an all-sky search, and 18 days for a directed search. For older, slower pulsars, on the other hand, we are only limited to 9 days for an all-sky search, and nearly 160 days for a directed search. The threshold sensitivities that these strategies can achieve, relative to the noise curves in Fig. 1, are plotted as functions of computing power in Fig. 2.

TABLE I. The number of independent parameter points $N_p(T, \mu_{\max} = 0.3)$ required for a coherent $T = 10^7$s search, for four fiducial types of pulsar. We list the requirements both for all-sky searches and for directed searches (i.e., searches where the source position is known in advance). Also listed are the threshhold values h_{th} of the characteristic strain h_c required to have 99% confidence of detection, assuming unlimited computer power. These threshold values are given by $h_{\mathrm{th}}/h_{3/\mathrm{yr}} = (1/1.90)\sqrt{\ln(50NN_p) - 1}$ where $N \equiv 2f_{\max}T$. Here $h_{3/\mathrm{yr}}$ is the corresponding threshhold, assuming the pulsar's postion and period and are known in advance.

Search Parameters		N_p	$h_{\mathrm{th}}/h_{3/\mathrm{yr}}$	N_p	$h_{\mathrm{th}}/h_{3/\mathrm{yr}}$
f (Hz)	τ (Yrs)	(all-sky)	(all-sky)	(directed)	(directed)
< 200	$> 10^3$	1.1×10^{10}	3.7	3.7×10^6	3.3
$< 10^3$	$> 10^3$	1.3×10^{16}	4.2	1.2×10^8	3.5
< 200	> 40	1.7×10^{18}	4.3	8.5×10^{12}	3.9
$< 10^3$	> 40	8×10^{21}	4.6	1.4×10^{15}	4.1

[2]This differs from equation (112) in [7] because we have specified 99% confindence, and we have use the correct exponential probability function for power.

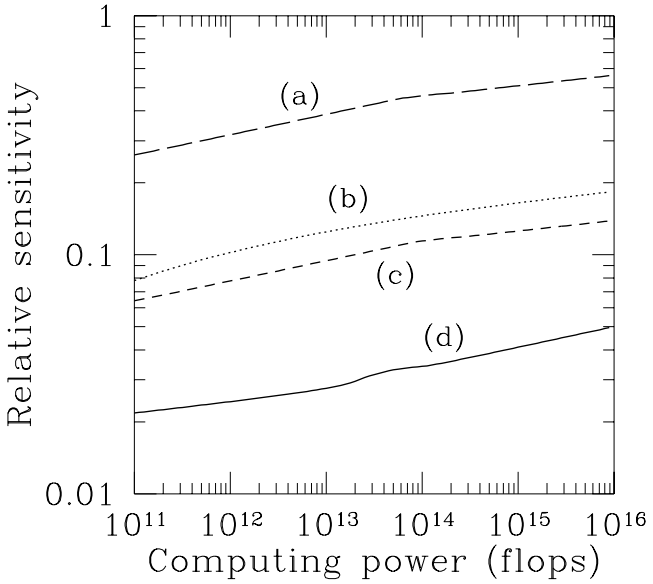

FIG. 2. Relative amplitude sensitivities $h_{3/\mathrm{yr}}/h_{1/\mathrm{T}}$ achievable with given computational resources, for various coherent search strategies: (a) directed search for old ($\tau \geq 1000$Yrs), slow ($f \leq 200$Hz) pulsars, (b) all-sky search for old, slow pulsars, (c) directed search for young ($\tau \geq 40$Yrs) fast ($f \leq 1000$Hz) pulsars, and, (d) all-sky search for these same sources. $h_{1/\mathrm{T}}$ is the expected sensitivity of the detector for an observation time T seconds, and with 99% confidence, assuming only that the frequency bandwidth of the source is constrained in advance. For a given computational power, we have determined the optimum observation time as described in sections VIB and VII.

E. Organization of this paper

In section II we outline the physics of pulsars which is relevant to the detection of continuous gravitational waves. The discussion is phenomenological and based almost entirely on pulsar data collected by radio astronomers. We focus attention on effects which may lead to significant frequency evolution over periods of several weeks of observation.

Then, in section III, we introduce a parameterised model of the expected gravitational waveform, including modulating effects due to detector motion.

From this, we go on in section IV to describe the basic technique used to search for signals, by constructing a demodulated time series. Livas [19], Jones [20] and Niebauer [21] have implemented variants of this basic search strategy over limited regions of parameter space (in particular they have not considered pulsar spin-down, and have restricted attention to small areas of the sky).

For the more computationally-intensive search over all sky positions and spindown parameters, it is important to be able to calculate the smallest number of independent parameter values which must be sampled in order to cover the entire space of signals. In section V we develop the metric formalism for calculating the number of independent points in parameter space.

In sections VI and VII we apply this formalism to determine the computational requirements of an all-sky search for unknown pulsars and a directed search, respectively.

Finally in section VIII, we list some possible alternatives to a straightforward coherent search of the interferometer data. Detailed studies of the pros and cons of each are currently under investigation.

II. PULSAR PHENOMENOLOGY

Currently, the only expected sources of continuous, periodic gravitational waves in the LIGO band are pulsars. In this section, therefore, we review those properties of pulsars which may be important in the detection process. In general, the search technique we present later is capable of detecting any *nearly* monochromatic gravitational wave with sufficient amplitude. However, it is useful to have a concrete physical system in mind when considering the expected gravitational waveform.

That pulsars are rapidly rotating neutron stars is now well established [22]. Their high densities and strong gravitational fields allow them to withstand rotation rates of hundreds of times per second. Moreover, pulsar emission mechanisms require large magnetic fields, frozen into (co-rotating with) the neutron star. Indeed these large field strengths may produce non-axisymmetric deformations of the pulsar. However, the most remarkable feature of pulsars is the very precise periodicity of observed pulses.

There are more than 700 known pulsars, all at galactic distances, concentrated in the galactic plane. Based on the sensitivity limits of radio observations the total number of active pulsars in our galaxy is estimated to be more than 10^5 [23,24].

A. Spindown

Pulsars lose rotational energy by electromagnetic braking, the emission of particles and, of course, emission of gravitational waves [25,26]. Thus, the rotational frequency is not completely stable, but varies over a timescale τ which is of order the age of the pulsar. Typically, younger pulsars (with periods of tens of milliseconds) have the largest spindown rates. Figure 3 shows the distribution of rotational frequencies and spindown age, $\tau = f/(df/dt)$.

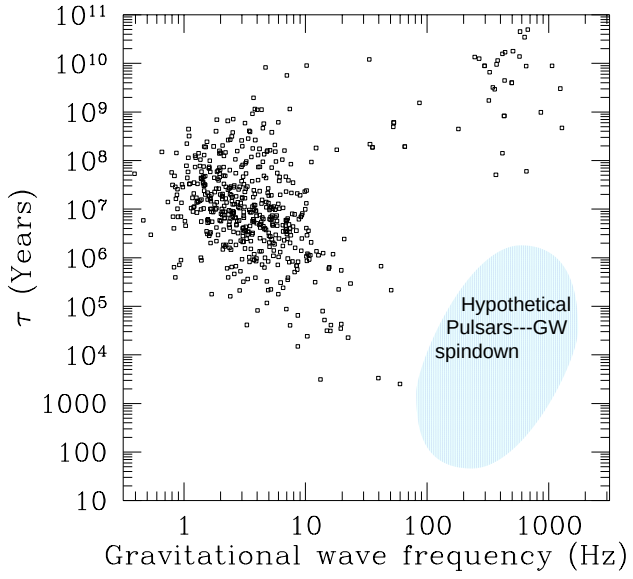

FIG. 3. Gravitational wave frequency versus spindown age, $\tau = f/(df/dt)$, measured in years, for 540 pulsars which have measured period derivative. The figure clearly shows a large concentration of pulsars in the mid-left of diagram. Most of these are isolated pulsars. The standard evolutionary scenario suggests that pulsars move from higher frequencies and shorter spindowns left and up towards this main bunch. In contrast, many of the millisecond pulsars lying in the upper right of the figure are in binary systems, and it is widely believed that these are pulsars which have been spun up by mass accretion from the companion star.

Current observations suggest that spindown is primarily due to electromagnetic braking; however, for detection purposes it is necessary to construct a sufficiently general model of the frequency evolution to cover all possibilities. For observing times $t_{\rm obs}$, say, much less than τ, the frequency drift is small and the rotational frequency[3] can be modeled as a power series of the form

$$f(t) = (f_0/2)\big(1 + \sum_k f_k t^k\big) \ . \qquad (2.1)$$

If $\tau_{\rm min}$ is the shortest timescale over which the frequency is expected to change by a factor of order unity, the coefficients satisfy

$$|f_k| \lesssim \tau_{\rm min}^{-k} \ . \qquad (2.2)$$

Clearly, for an observation time $t_{\rm obs} \ll \tau_{\rm min}$, the first few terms in this series will dominate.

Observations suggest that pulsars are born in supernova explosions with very short periods (perhaps several milliseconds), and subsequently spin down on timescales comparable to their age. Supernovae are observed in galaxies similar to our own at the rate of two or three per century, so we might expect $\tau_{\rm min} \sim 40$ years for pulsars in our galaxy. It is at this point that the distinction between various classes of pulsars becomes important. The known millisecond pulsars are old neutron stars which have have been spun up to periods of only a few milliseconds, possibly by episodes of mass transfer from a companion star. As seen from Fig. 3, timing measurements of millisecond pulsars yield very long spindown timescales, $\tau_{\rm min} \gtrsim 10^7$ years.

B. Proper Motions

Pulsars are generally high velocity objects [25], as can be inferred by the distance they move in their lifetimes. Proper motions cause Doppler shifts in the observed pulsar frequency. If the motion is uniform (constant velocity), it simply induces a constant frequency shift — an effect which is undetectable. However, acceleration and higher order derivatives of the source's motion will modulate the observed frequency.

Studies of millisecond pulsars in globular clusters have shown that acceleration in the cluster field can produce frequency drifts which are comparable in magnitude to the spindown effects [27,28]. Once again, we expect these effects to be well modelled by a power series in $t/\tau_{\rm cross}$, where $\tau_{\rm cross}$ is the time it takes the pulsar to cross the cluster. We expect that $\tau_{\rm min} \leq \tau_{\rm cross}$ for these objects (since if not, the pulsar will already have escaped the cluster). Thus the frequecy model adopted above should be sufficiently general to encompass the observational effects of proper motions of the sources.

A large proportion of millisecond pulsars are also in binary systems. Unfortunately, such pulsars participate in proper motions which vary over very short timescales (their orbital periods). The time-dependent Doppler effect due to these motions is *not* modelled accurately by a simple power series as in Eq. (2.1). They would require a more elaborate model involving as many as five unknown orbital parameters. Including these effects in a coherent, all-sky search strategy would be prohibitive (see section VI). In a search for gravitational waves from a known binary pulsar, however, it would be important to deal with this effect.

Proper motions can also affect a search if the star moves across more than one resolution element on the sky during an observation. For the lengths of observation periods envisioned here, this is unlikely to be a problem. In an observation lasting a year, however, a pulsar with a spatial velocity of 1×10^3 km s^{-1} at a distance of 300 pc will move by about half an arcsecond, which is comparable to the resolution limit for our observations if the pulsar frequency is 1kHz.

[3]We choose to parameterise the frequency by what will be the gravitational wave frequency, f_0, thus introducing the extra factor of 2 into this expression

C. Glitches

In addition to gradual frequency drifts due to spin-down, some young pulsars exhibit occasional, abrupt increases in frequency. The physical mechanism behind these frequency *glitches* is not well understood, although the number of observations of glitch events is growing [23]. Given the stochastic nature of glitching, and the expectation that several months will elapse between major events, we will ignore glitching in this paper.

III. GRAVITATIONAL WAVES FROM PULSARS

In order to gain insight into the detection problem it is also important to understand the expected gravitational wave signal. Several mechanisms have been discussed in the literature which may produce non-axisymmetric deformations of a pulsar, and hence lead to gravitational wave generation [12–14,16,29,30].

In general, a pulsar can radiate strongly at frequencies other than twice the rotation frequency. For example, a pulsar deformed by internal magnetic stresses, which are not aligned with a principal axis, can radiate at the rotation frequency and twice that frequency [31]. If the star precesses, it will radiate at three frequencies: the rotation frequency, and the rotation frequency plus and minus the precession frequency [16]. The important point, however, is that the signal at the detector is generally narrow band, exhibiting only slow frequency drift on observational timescales.

Therefore, in this section we outline the main features of the expected waveform and the corresponding strain measured at a detector for the case of crustal deformation; other scenarios give similar results except for the presence of more than one spectral component.

A. Waveform

Adopting a simple model of a distorted pulsar as a tri-axial ellipsoid, rotating about a principal axis with a frequency given by Eq. (2.1), one may compute the expected gravitational wave signal using the quadrupole formula. The two polarizations are

$$h_+ = h_0(1+\cos^2 i)\ \cos\{2\pi\, f_0\,[t + \textstyle\sum f_k \frac{t^{k+1}}{k+1}]\}\ , \qquad (3.1)$$

$$h_\times = 2h_0 \cos i\ \sin\{2\pi\, f_0\,[t + \textstyle\sum f_k \frac{t^{k+1}}{k+1}]\}\ , \qquad (3.2)$$

where i is the angle between the rotation axis and the line of sight to the source. The dimensionless amplitude is

$$h_0 = \frac{8\pi^2 G}{c^4}\ \frac{I_{zz} f_0^2}{r}\ \epsilon\ , \qquad (3.3)$$

where

$$\epsilon = \frac{I_{xx} - I_{yy}}{I_{zz}} \qquad (3.4)$$

is the *gravitational ellipticity* of the pulsar. The distance to the source is r, and I_{jk} is its moment of inertia tensor.

The strength of potential sources is best discussed in terms of the characteristic amplitude h_c, defined in Eq. (50) of [7], and simply related to h_0 by

$$h_c = \sqrt{\frac{2}{15}} h_0\ . \qquad (3.5)$$

For a typical $1.4 M_\odot$ neutron star, having a radius of 10km and at a distance of 10kpc, the dimensionless amplitude is

$$h_c = 7.7\times 10^{-25}\ \frac{\epsilon}{10^{-5}}\ \frac{I_{zz}}{10^{45}\ \mathrm{g\ cm^2}} \frac{10\mathrm{kpc}}{r} \left(\frac{f_0}{1\mathrm{kHz}}\right)^2\ . \qquad (3.6)$$

The magnitude of the gravitational ellipticity, ϵ, represents the central uncertainty in any estimate of gravitaional waves from pulsars. The tightest theoretical constraint, $\epsilon < 10^{-5}$, is set by the maximum strain that the neutron star crust may support [7]. It has also been suggested that stresses induced by large magnetic fields might result in significant gravitational ellipticity. Recently, Bonazzola and Gourgoulhon [12] have considered this possibility, finding discouraging results; their calculations indicate $10^{-13} \lesssim \epsilon \lesssim 10^{-9}$ depending on the precise model they consider.

In any case, an upper bound on the gravitational ellipticity is $\epsilon \sim 10^{-5}$, although typical values may be significantly smaller.

B. Signal at the detector

Observing the gravitational waves using an earth-based interferometer introduces two further difficulties into the detection process: Doppler modulation of the observed gravitational wave frequency, and amplitude modulation due to the changing orientation of the detector.

For the purpose of detection, the Doppler modulation of the observed gravitational wave frequency, due to motion of the detector with respect to the solar system barycenter, is a large effect. Assuming the intrinsic frequency model (2.1) for the pulsar rotation, the gravitational wave frequency measured at the detector is

$$f_{\mathrm{gw}}(t) = f_0\left(1 + \sum_k f_k t^k\right)\left(1 + \frac{\vec{v}}{c}\cdot\hat{n}\right) \qquad (3.7)$$

where $\vec{v}(t)$ is the detector velocity and $\hat{n}$ is the unit vector pointing to the pulsar, in some inertial frame. We generally choose this frame to be initially comoving with the Earth at $t = 0$. The frequency measured in this frame is

identical to that measured at the solar system barycenter except for an unimportant constant shift in f_0.

To understand the amplitude modulation we must introduce the Euler angles, $\{\Theta, \Phi, \Psi\}$, which specify the orientation of the gravitational wave frame with respect to the detector frame. The dimensionless strain at the detector is

$$h = F_+(\Theta, \Phi, \Psi)h_+ + F_\times(\Theta, \Phi, \Psi)h_\times \tag{3.8}$$

where F_+ and $F_\times$ are the detector beam patterns given in Thorne [7]. In searching for continuous gravitational waves from a particular direction, the Euler angles become periodic function of sidereal time, thus resulting in an amplitude and phase modulation of the observed signal [7,12,19]. For observation times longer than one sidereal day, the amplitude modulation effectively averages the reception over all values of right ascension, and over a range of declination which depends one the precise position of the pulsar. In particular, the effect of this process is to allow detection of continuous waves from any direction, but at the cost of reducing the measured strain (see Fig. 4).

C. Parameter space

To facilitate later discussion it is useful to parameterize the gravitational waveform by a vector $\boldsymbol{\lambda} = (\lambda^0, \vec{\lambda})$ such that

$$(\lambda^0, \lambda^1, \ldots \lambda^{s+2}) = (f_0, n_x, n_y, f_1, \ldots, f_s) \; . \tag{3.9}$$

Here s is the maximum number of spindown parameters included in the frequency model determined by Eq. (2.1). These vectors span an $s+3$ dimensional space on which λ^α can be thought of as coordinates. (Note that $n_z^2 = 1 - n_x^2 - n_y^2$ is not an independent parameter.) In particular we denote the observed phase of the gravitational waveform by

$$\phi(t; \boldsymbol{\lambda}) = 2\pi \int^t dt' \; f_{\rm gw}(t') \; , \tag{3.10}$$

where $f_{\rm gw}(t')$ is given by Eq. (3.7).

Initial interferometers in LIGO should have reasonable sensitivity to gravitational waves with frequencies

$$f \geq 40{\rm Hz} \; , \tag{3.11}$$

while advanced interferometers are expected to have improved sensitiviy down to

$$f \geq 10{\rm Hz} \; . \tag{3.12}$$

Moreover, theoretical constraints suggest that pulsars with spin periods significantly smaller than one millisecond are unlikely. This helps to constrain the highest frequency that one may wish to consider in an all sky search to be about 2kHz. According to the discussion in section II, the spindown parameters satisfy

$$-\tau_{\rm min}^{-k} \leq f_k \leq \tau_{\rm min}^{-k} \; , \tag{3.13}$$

where $\tau_{\rm min}$ is the minimum spindown age of a pulsar to be searched for. Finally, n_x and n_y are restricted by the relation

$$n_x^2 + n_y^2 \leq 1 \; . \tag{3.14}$$

IV. DATA ANALYSIS TECHNIQUE

Radio astronomers are familiar with searching for nearly periodic sources in the output of their detectors [28,32]. The technique employed by them is directly applicable to the problem at hand [19,20].

In the detector frame the gravitational wave signal can be written as

$$h(t; \boldsymbol{\lambda}) = {\rm Re}\left[\mathcal{A} e^{-i\phi(t;\boldsymbol{\lambda})}\right] \tag{4.1}$$

where $\mathcal{A} = (h_{0+} + ih_{0\times})$, $h_{0+} = F_+(1+\cos^2 i)h_0$ and $h_{0\times} = 2F_\times(\cos i)\, h_0$. The orbital phase $\phi(t; \boldsymbol{\lambda})$ is given by Eqs. (3.10) and (3.7). Introducing a canonical time

$$t_b[t; \vec{\lambda}] = \frac{\phi(t; \boldsymbol{\lambda})}{2\pi f_0} \; , \tag{4.2}$$

the above signal becomes monochromatic as a function of t_b. (The presence of the amplitude modulation complicates the following analysis without changing the conclusions significantly; therefore, we treat $\mathcal{A}$ as constant in this and the next section[4].) Figure 4 shows the normalised power spectrum computed from the signal as a function of t in Eq. (4.1) (with $f_k \equiv 0$), compared with the spectrum from the signal as a function of t_b. It is clear that the maximum power per frequency bin is significantly reduced when frequency modulation is not accounted for.

[4] Amplitude modulation can be viewed as the convolution of the exactly periodic signal with some complicated window function. Thus, in reality, the power spectrum of a stretched signal will not be a monochromatic spike at a single frequency, but will be split into several discrete, narrow spikes spread over a bandwidth $\delta f \simeq 10^{-4}$Hz. After a preliminary detection, the amplitude modulation spikes would provide a discriminant against false signals [19].

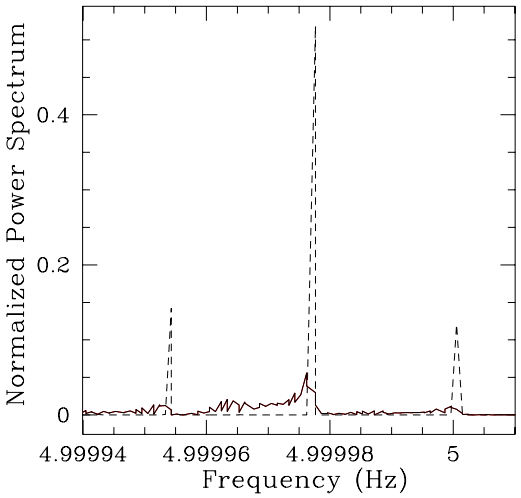

FIG. 4. Power spectra for two signals, each with gravitational wave frequency 5Hz, computed using approximately 10 days worth of data; they are normalised with respect to the maximum power achieved if the source was directly above an interferometer which remained stationary during the entire observation. The signal was assumed to come from declination 0° and right ascension 90°; in fact the amplitude modulation is only sensitive to changes in declination. The detector latitude was chosen to coincide with LIGO detector in Hanford Washington. The solid line corresponds to a Doppler and amplitude modulated gravitational wave signal. The dashed line is the same signal but with the Doppler modulation removed by stretching. The (unreasonably) low frequency was chosen for illustrative purposes, so that both curves could appear on the same scale. For realistic gravitational wave frequencies ($\sim$ 500Hz) the Doppler modulated signal would be further reduced by roughly two orders of magnitude.

Radio astronomers refer to this technique of introducing a canonical time coordinate as *stretching* the data. Since interferometer output will be sampled at approximately 16kHz, in a practical search for pulsars up to 2kHz gravitational wave frequency, the stretching can probably be achieved by resampling the data stream appropriately. This method, which is called *stroboscopic sampling* by Schutz [18], has the benefit of keeping the computational overhead introduced by the stretching process to a minimum. We will return to this issue in a later publication.

Now, a search of the detector output, $o(t)$, for gravitational waves from a known source is straightforward. One assumes specific parameter values $\vec{\xi}$ in the waveform (4.1), computes the demodulated time function $t_b[t;\vec{\xi}]$ using Eq. (4.2) and stretches the detector output accordingly, thus

$$o_b(t_b[t;\vec{\xi}]) = o(t)\ . \tag{4.3}$$

If the assumed parameters $\vec{\xi}$ are not too much different from the actual parameters $\vec{\lambda}$ of the signal, the stretched data will consist of a nearly monochromatic signal. One then takes the Fourier transform with respect to t_b,

$$\tilde{o}(f;\vec{\xi}) = \frac{1}{\sqrt{T_b^{\rm obs}}}\int_0^{T_b^{\rm obs}} e^{2\pi i f t_b}\, o_b(t_b)\, dt_b\ . \tag{4.4}$$

Here $T_b^{\rm obs}$ is length of the observation measured using t_b. The power spectrum is then searched for excess power. (The threshold is set by demanding some overall statistical significance for a detection; see section VI.) Notice that the gravitational wave frequency, $\lambda^0 = f_0$, is treated somewhat differently than the other parameters; the Fourier transform searches over all possible values in a single pass. Given a sampled data set containing N points, the entire process, from original data through to the power spectrum, requires of order $3N\log_2 N$ floating point operations (to first approximation).

If all the parameters are not known accurately in advance, it will be necessary to search over some of the remaining parameters $\vec{\lambda}$; a separate demodulation and FFT must be performed for each independent point in parameter space that one wishes to search. There are many possible refinements on this strategy which could reduce the computational cost of a search by circumventing certain stages of the procedure described here. We mention some of them in section VIII, however, we focus attention on this baseline strategy in this paper.

One more issue that arises in the discussion of stretching is how it effects the noise in the detector. Throughout this paper we assume that the noise in the detector is a stationary, gaussian process; however, when we stretch the output data stream the noise is no longer strictly stationary unless it is perfectly white. Real detectors will have coloured noise, with correlations between points sampled at different times. Stretching the data modifies these correlations in a time dependent manner. In our case this is a very small effect, having a characteristic timescale of several hours, and besides this the noise in real detectors may be intrinsically non-stationary on similar timescales due to instrumental effects. Correcting pulsar searches for such non-stationarity is an important problem, but one that we do not address here. We simply assume that $S_n(f)$, the power spectral density of the noise, can be estimated on short timescales and used in the conventional way for signal to noise estimates. Moreover, the effects of stretching on noise are only a consideration when the noise is not white; since stretching affects the power spectrum only within bands $\sim 10^{-1}$Hz wide, the detector spectrum can usually be taken as white, unless we are near a strong feature in the noise spectrum. The precise nature of these effects is being explored by Tinto [33].

V. PARAMETER SPACE METRIC

In general, neither the position of the pulsar nor its intrinsic spindown may be known in advance of detection. Therefore, the above process, or some variant on it,

must be repeated for many different vectors $\vec{\xi}$ until the entire parameter space has been explored. How finely must one sample these parameters in order to minimise the risk of missing a signal? A similar question arises in the context of searching for signals from coalescing compact binaries using matched filtering; Owen [5] has introduced a general framework to provide an answer in that case. We adapt his method to the problem at hand by defining a distance function on our parameter space; the square of distance between two points in parameter space is proportional to the fractional loss in signal power due to imprecise matching of parameters. The number of discrete points which must be sampled can then be determined from the proper volume of the parameter space with respect to this metric.

A. Mismatch

The one-sided power spectral density (PSD) of the detector output, stretched with parameters $\vec{\xi}$, is

$$P_o(f) = 2\,|\tilde{o}(f;\vec{\xi})|^2 \ . \tag{5.1}$$

Now, suppose a detector output consists of a signal with parameters $\boldsymbol{\lambda}$, and stationary, gaussian noise $n(t)$ such that

$$o(t) = h(t;\boldsymbol{\lambda}) + n(t) \ . \tag{5.2}$$

Thus, the expected PSD of the detector output, once again stretched with parameters $\vec{\xi}$, is

$$E[\,P_o(f)\,] = 2\,|\tilde{h}(f;\boldsymbol{\lambda},\Delta\vec{\lambda})|^2 + S_n(f) \ , \tag{5.3}$$

where $\Delta\vec{\lambda} = \vec{\xi} - \vec{\lambda}$, and $S_n(f)$ is the one-sided power spectral density of the detector noise. (As discussed at the end of the previous section, we ignore the small effects of stretching on the noise.) The notation $\tilde{h}(f;\boldsymbol{\lambda},\Delta\vec{\lambda})$ indicates the Fourier transform of a signal, with parameters $\boldsymbol{\lambda}$, with respect to a time coordinate $t_b[t;\vec{\lambda}+\Delta\vec{\lambda}]$. We define the *mismatch* $m(\boldsymbol{\lambda},\Delta\boldsymbol{\lambda})$ to be the fractional reduction in signal power caused by stretching the data with the wrong parameters, *and* by sampling the spectrum at the wrong frequency; specifically,

$$m(\boldsymbol{\lambda},\Delta\boldsymbol{\lambda}) = 1 - \frac{|\tilde{h}(f;\boldsymbol{\lambda},\Delta\vec{\lambda})|^2}{|\tilde{h}(f_0;\boldsymbol{\lambda},0)|^2} \ . \tag{5.4}$$

Remember that $\boldsymbol{\lambda} = (\lambda^0 = f_0, \vec{\lambda})$.

In the present circumstance, it is sufficient to consider a complex signal

$$h(t;\boldsymbol{\lambda}) = \mathcal{A} e^{-2\pi i f_0 t_b[t;\vec{\lambda}]} \ , \tag{5.5}$$

where the amplitude $\mathcal{A}$ is constant. The function $t_b[t;\vec{\lambda}]$, computed using Eqs. (4.2), (3.10) and (3.7), is explicitly written as

$$t_b[t;\vec{\lambda}] = \int_0^t dt' \{(1+\textstyle\sum f_k t'^k)(1+\vec{v}\cdot\hat{n}/c)\} \ , \tag{5.6}$$

Now, the Fourier transform $\tilde{h}(f;\boldsymbol{\lambda},\Delta\vec{\lambda})$ is

$$\tilde{h}(f;\boldsymbol{\lambda},\Delta\vec{\lambda}) = \frac{\mathcal{A}}{\sqrt{T_b^{\rm obs}}} \int_0^{T_b^{\rm obs}} d\hat{t}_b \; e^{i\Phi[t;\boldsymbol{\lambda},\Delta\vec{\lambda}]} \ , \tag{5.7}$$

where

$$\frac{\Phi[t;\boldsymbol{\lambda},\Delta\vec{\lambda}]}{2\pi} = \Delta\lambda^0 \hat{t}_b + f_0(t_b[t;\vec{\lambda}+\Delta\vec{\lambda}] - t_b[t;\vec{\lambda}]) \tag{5.8}$$

and $\Delta\lambda^0 = f - f_0$. Here, t should be interpreted as a function of $\hat{t}_b$ defined implicity by $\hat{t}_b = t_b[t;\vec{\lambda}+\Delta\vec{\lambda}]$. Using Eqs. (5.6)-(5.8) it is easy to show that $m(\boldsymbol{\lambda},\Delta\boldsymbol{\lambda})$ has a local minimum of zero when $\Delta\boldsymbol{\lambda} \equiv 0$;

$$m(\boldsymbol{\lambda},\Delta\boldsymbol{\lambda})|_{\Delta\boldsymbol{\lambda}=0} = 0 \ , \tag{5.9}$$

$$\partial_{\Delta\lambda^\alpha} m(\boldsymbol{\lambda},\Delta\boldsymbol{\lambda})|_{\Delta\boldsymbol{\lambda}=0} = 0 \ . \tag{5.10}$$

Thus, an expansion of the mismatch in powers of $\Delta\boldsymbol{\lambda}$ is

$$m(\boldsymbol{\lambda},\Delta\boldsymbol{\lambda}) = \sum_{\alpha,\beta} g_{\alpha\beta}(\boldsymbol{\lambda})\Delta\lambda^\alpha\Delta\lambda^\beta + \mathcal{O}(\Delta\boldsymbol{\lambda}^3) \ , \tag{5.11}$$

where

$$g_{\alpha\beta} = \tfrac{1}{2}\sum_{\alpha,\beta} \partial_{\Delta\lambda^\alpha}\partial_{\Delta\lambda^\beta} m(\boldsymbol{\lambda},\Delta\boldsymbol{\lambda})|_{\Delta\boldsymbol{\lambda}=0} \ . \tag{5.12}$$

In this way the mismatch defines a local distance function on the signal parameter space, and, for small separations $\Delta\boldsymbol{\lambda}$, $g_{\alpha\beta}$ is the metric of that distance function. Note that the metric formulation (5.11) will generally *overestimate* the mismatch for large separations, as demonstrated in Figure 5.

Calculations using this formalism are considerably simplified by partially evaluating the right hand side of Eq. (5.12). The form of the signal (5.5) allows us to write

$$g_{\alpha\beta}(\boldsymbol{\lambda}) = \langle \partial_{\Delta\lambda^\alpha}\Phi\partial_{\Delta\lambda^\beta}\Phi\rangle - \langle\partial_{\Delta\lambda^\alpha}\Phi\rangle\langle\partial_{\Delta\lambda^\beta}\Phi\rangle \ , \tag{5.13}$$

where Φ is given by Eq. (5.8), and where we use the notation

$$\langle\ldots\rangle = \frac{1}{T_b^{\rm obs}} \int_0^{T_b^{\rm obs}} (\ldots) dt_b \Big|_{\Delta\boldsymbol{\lambda}=0} \ . \tag{5.14}$$

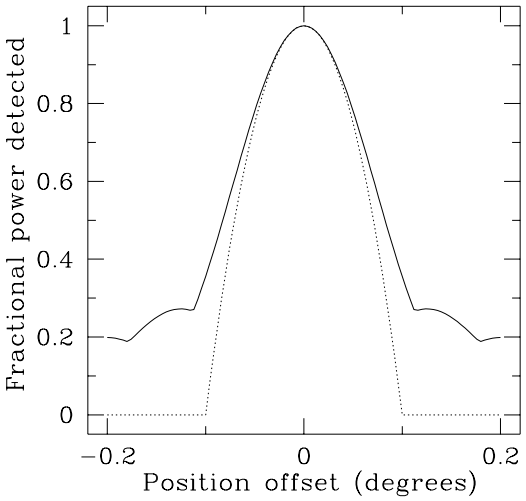

FIG. 5. Fractional reduction in measured signal power caused by demodulating by mismatched parameters (in this case, an error in the assumed declination of the source). The solid curve is the true power ratio, the dotted is that given by the quadratic approximation of the metric. Note that the widths of the curves agree well down to 70% power reduction ($m \sim 0.7$), beyond which the metric approximation significantly underestimates the range of parameters permitted for a specified power loss. The curves are computed for a sky position of 0° right ascension, 45° declination, and no spindown.

B. Metric and number of patches

Up until now, we have treated the frequency of the signal as one of the parameters, λ_0, which must be matched. In our search technique, stretching and Fourier transforming the data yields an entire power spectrum, automatically sampling all possible frequencies. We would really like to know the number of times that this combination of procedures must be performed in a search. This requires knowledge of the mismatch $m(\boldsymbol{\lambda}, \Delta\boldsymbol{\lambda})$ as a function of $\Delta\vec{\lambda}$, having already maximized the power (i.e. minimized m) over λ^0. The result is the mismatch projected onto the $(s+2)$-parameter subspace:

$$\mu = \min_{\lambda_0} m(\boldsymbol{\lambda}, \Delta\boldsymbol{\lambda}) = \sum_{ij} \gamma_{ij} \Delta\lambda^i \Delta\lambda^j \tag{5.15}$$

where

$$\gamma_{ij} = g_{ij} - \frac{g_{0i} g_{0j}}{g_{00}} \,, \tag{5.16}$$

and $i = 1, \ldots, s+2$. We will generally refer to μ as the *projected mismatch*.

Technically, γ_{ij} should be computed from $g_{\alpha\beta}$ evaluated at the specific value of λ_0 at which the minimum projected mismatch occurred. However, since this number is unknown in advance of detection, we evaluate γ_{ij} for the largest frequency in the search space. In this way we never underestimate the projected mismatch.

In a search, the parameter space will be sampled at a lattice of points, chosen so that no location in the space has μ (given by Eq. (5.15) greater than some $\mu_{\max}$ away from one of the points. This is equivalent to tiling the parameter space with *patches* of maximum extent $\mu_{\max}^{1/2}$. The number of points we must sample at is therefore

$$N_p = \frac{\int_{\mathcal{P}} \sqrt{\det ||\gamma_{ij}||}\, d^{s+2}\vec{\lambda}}{V_{\text{patch}}} \,, \tag{5.17}$$

where V_{patch} is the proper volume of a single patch, and $s+2$ is the reduced dimensionality of the parameter space $\mathcal{P}$ (excluding λ_0).

Optimally, one should use some form of spherical closest packing to cover the space with the fewest patches. Our solution uses hexagonal packing in two of the dimensions and cubic packing in all the others; in this way the volume of a single patch is

$$V_{\text{patch}} = \frac{3\sqrt{3}}{4} \left(\frac{4\mu_{\max}}{s+2} \right)^{(s+2)/2} \tag{5.18}$$

VI. DEPTH OF AN ALL SKY SEARCH

We are finally in a position to estimate the depth of a search for periodic sources using LIGO. The detector participates in two principal motions which cause significant Doppler modulations of the observed signal: daily rotation, and revolution of the Earth about the Sun. The latter is actually a complex superposition of an elliptical Keplerian orbit with a smaller orbit about the earth-moon barycenter, and is further perturbed by interactions with other planets. For now, however, we use a simplified model which treats both rotation and revolution as circular motions about separate axes inclined at an angle $\epsilon = 23°27'$ to each other. Although a simplification, this does remove any spurious symmetries from the model; thus, an actual search using the precise ephemeris of the earth in its demodulations should give comparable results. In this model, then, we write the velocity of the detector in a frame which is inertial to the solar system barycenter but initially comoving with the earth:

$$\begin{aligned} \vec{v} = &-(\Omega R_d \sin\Omega t - \Omega_A R_A \sin\Omega_A t)\vec{x} \\ &+(\Omega R_d \cos\Omega t - \Omega_A R_A \cos\epsilon[\cos\Omega_A t - 1])\vec{y} \\ &-\Omega_A R_A \sin\epsilon[\cos\Omega_A t - 1])\vec{z} \end{aligned} \tag{6.1}$$

where $R_d = 6.371 \times 10^8 (\cos l)$cm, l is the latitude of the detector, and $R_A = 1.496 \times 10^{13}$cm is the distance from the earth to the sun. The angular velocities are $\Omega = 2\pi/(86\,400\text{s})$ and $\Omega_A = 2\pi/(3.155674 \times 10^7\text{s})$. Our coordinate system measures $\vec{x}$ towards the vernal equinox and $\vec{z}$ towards the north celestial pole, and we arbitrarily choose to measure time starting at noon on the vernal equinox.

The number of spindown parameters f_k which must be included to account for all intrinsic frequency drift depends to a large extent on the type of pulsar one wishes to search for. We determined this number on a case by case basis, including all parameters which lead to a significant increase in the number of parameter space patches. Equivalently, the following geometric picture suggests a simple criterion for deciding when there is one spindown parameter too many included in the signal parametrization. Let λ^L be the last, 'questionable' spindown parameter f_s (so $L = s+2$). With respect to the natural metric γ_{ij} on parameter space, the unit-normal to surfaces of constant λ^L is just $\gamma^{iL}/(\gamma^{LL})^{1/2}$, where γ^{ij} is the inverse of γ_{ij}. The spindown parameter λ^L is unnecessary if the proper thickness of the parameter space in this normal direction is less than half the proper grid spacing; that is, if $2\tau_{\rm min}^{-L+2}/(\gamma^{LL})^{1/2} < \sqrt{\mu_{\rm max}/L}$. In practice, one has included more spindown parameters than necessary if and only if $\gamma^{LL} > 4L\tau_{\rm min}^{-2L+2}/\mu_{\rm max}$.

A. Patch number versus observation time

It is extremely difficult to obtain a closed form expression for the metric, let alone its determinant. Therefore, we present results for two concrete scenarios which suggest themselves based on the discussion in section II: (i) hypothetical sources with $f_0 \leq 1000$Hz, and spindown ages greater than $\tau = 40$Yrs; incidentally, this also includes the majority of known, millisecond pulsars; and (ii) slower sources ($f_0 \leq 200$Hz) having spindown ages in excess of $\tau = 1000$Yrs. The number of parameter space points which must be searched is plotted as a function of total observation time in Fig. 6. The numbers are normalised by a maximum projected mismatch $\mu_{\rm max} = 0.3$.

In considering an optimal choice of observation time, it is useful to construct an empirical fit to $N_p(t_{\rm obs}, \mu_{\rm max})$. Notice first that all the parameters $\Delta\vec{\lambda}$ in Φ, given by Eq. (5.8), appear multiplied by the gravitational wave frequency f_0; thus, $N_p \propto (f_0)^{s+2}$. Furthermore, provided the determinant of the metric is only weakly dependent on the values of the f_k one may also extract a factor of $\tau^{-s(s+2)/2}$; our investigations suggest the validity of this approach. In this way we arrive at the expression

$$N_p \simeq \max_{s\in\{0,1\ldots\}} \left[\mathcal{N}_s F_s(t)\right] \ , \tag{6.2}$$

where

$$\mathcal{N}_s = \left(\frac{f_0}{1\text{kHz}}\right)^{s+2} \left(\frac{40\text{Yrs}}{\tau}\right)^{s(s+1)/2} \left(\frac{0.3}{\mu_{\rm max}}\right)^{(s+2)/2} \tag{6.3}$$

$$F_0(t_{\rm obs}) = 6.9\times 10^3\, T^2 + 3.0\, T^5 \tag{6.4}$$

$$F_1(t_{\rm obs}) = \frac{1.9\times 10^8\, T^8 + 5.0\times 10^4\, T^{11}}{4.7 + T^6} \tag{6.5}$$

$$F_2(t_{\rm obs}) = \frac{2.2\times 10^7\, T^{14}}{56.0 + T^9} \ , \tag{6.6}$$

and T is the observation time measured in days. These formulae are normalised using only the data corresponding to Fig 6(a), and subsequently compared with computed values for several frequencies and spindown ages τ. The analytic fit is in good agreement with the computed results for a variety of parameters; however, the fits generally break down for observation times less than one day. We stress that more spindown parameters may become important for observation times longer than 30 days.

Schutz [18] has previously estimated the number of points which must be searched in the abscence of spindown corrections; he argued that this number scaled as T^4 for observation times longer than about a day. The difference between his previous estimate and the expression in Eq. (6.4), which shows that the number of points increases as T^5, derives from an asymmetery between declination and right ascension which was not accounted for in his argument.

The benefit of the metric formulation is that it accounts for the significant correlations which exist between the intrinsic spindown and the earth-motion-induced Doppler modulations by using points which lie on the principal axes of the ellipsoids described by Eq. (5.15). Replacing the invariant volume integral in Eq. (5.17) by

$$\int_{\mathcal{P}} \sqrt{\textstyle\prod_i \gamma_{ii}}\; d^{s+2}\vec{\lambda} \ , \tag{6.7}$$

gives the number of points required for a search if, instead, one chooses them to lie on the $\{n_x, n_y, f_1, f_2, ...\}$ coordinate grid. Figure 7 shows the total number of points computed using this method compared to the results obtained using the invariant volume integral. For sufficiently long integration times the difference can be several orders of magnitude.

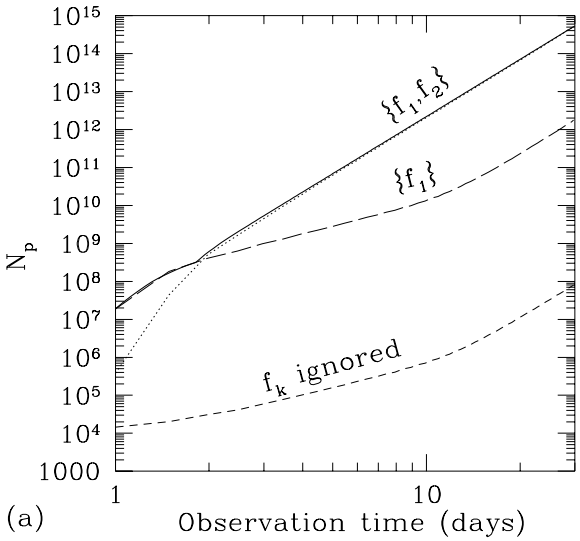

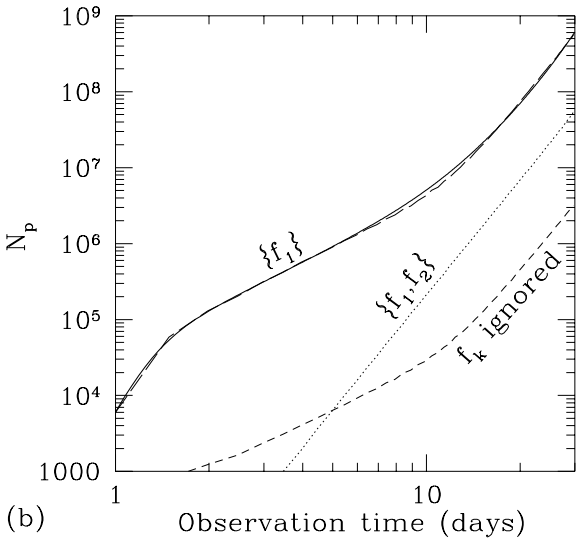

FIG. 6. Number of independent points in parameter space as a function of total observation time, using a maximum projected mismatch $\mu_{\max} = 0.3$. The parameter ranges chosen were: (a) maximum gravitational wave frequency 1 000Hz, minimum spindown age $\tau_{\min} = 40$Yrs (hypothetical young pulsars); (b) maximum gravitational wave frequency 200Hz, minimum spindown age $\tau_{\min} = 10^3$Yrs (observed, slow pulsars). The short-dashed curve represents the total number of patches ignoring all f_k. The long-dashed curve is the number of patches including only f_1 in the search. The dotted line is the number of patches including both f_1 and f_2. Also shown is the empirical fit given in the text; it was normalised by the results in shown in (a). In some regimes, searching over an additional spindown parameter would seem to reduce the number of patches; however, this actually only indicates regions where the parameter space extends less than one full patch width in the additional dimension. In such regimes one must properly discard the extra parameter from the search, forcing one to choose always the higher of the curves.

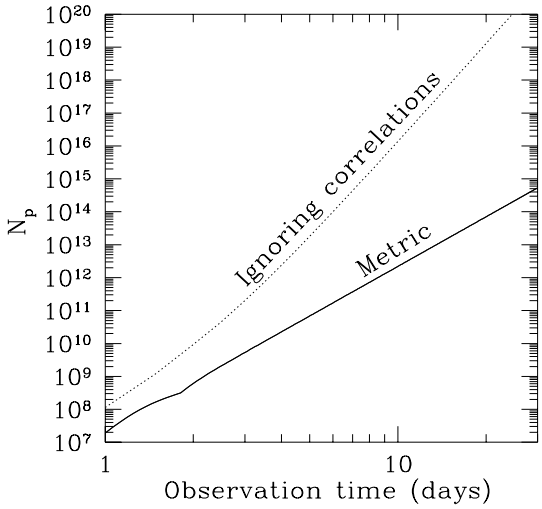

FIG. 7. The total number of parameter-space points needed to search for pulsars having gravitational wave frequency up to 1kHz, and spindown age greater than $\tau = 40$Yrs. The solid line is the number computed using the metric and properly accounting for correlations between various terms in the frequency evolution. The dotted line is the same number computed directly by assuming the points must lie on the grid of coordinates used to parameterise the signal. The benefits of using the metric to optimally place the points to be searched in parameter space is clear.

B. Computational Requirements

The number of real samples of the interferometer output for an observation lasting $t_{\rm obs}$ seconds, and sampled at a frequency $2f_{\max}$, where $f_{\max}$ is the maximum gravitational wave frequency being searched for, is

$$N = 2f_{\max}t_{\rm obs} \; . \tag{6.8}$$

For each $\vec{\lambda}$ that is used to stretch the detector output, a search then requires an FFT, calculation of the power, and some thresholding test for excess power. Assuming that the stretching and thresholding require negligible computations compared to performing the FFT and computing the power, the total number of floating point operations for a search is

$$N_{\rm op} = 6f_{\max}t_{\rm obs}N_p[\log_2(2f_{\max}t_{\rm obs}) + 1/2] \; , \tag{6.9}$$

where N_p is given by (6.2)-(6.6). The additive 1/2 inside the square brackets accounts for the three floating point operations per frequency bin which is required to compute the power from the Fourier transform.

A guideline for a feasible, long-term, search strategy is that data reduction should proceed at a rate comparable to data acquisition. Thus, the total computing power required for data reduction, in floating point operations per second (flops), is

$$P = \frac{N_{\rm op}}{t_{\rm obs}} = 6 f_{\rm max} N_p(t_{\rm obs}, \mu_{\rm max})[\log_2(2 f_{\rm max} t_{\rm obs}) + 1/2] \ . \tag{6.10}$$

For a prescribed maximum projected mismatch $\mu_{\rm max}$, and maximum available computing power $P_{\rm max}$ this expression determines the maximum allowed coherent integration time. Alternatively, given the computing power available for data reduction, $P_{\rm max}$, it provides an implicit relation between $\mu_{\rm max}$ and the integration time.

The idea now is to choose $\mu_{\rm max}$ and $t_{\rm obs}$ so that we maximize the sensitivity of the search. In order to do this we must first obtain a threshold, above which we consider excess power to indicate the presence of a signal.

As discussed in section IV, we assume that the noise in the detector is a stationary, gaussian random process with zero mean and PSD $S_n(f)$. In the abscence of a signal, the power $P_o(f) = 2|\tilde{n}(f)|^2$ is exponentially distributed with probability density function

$$\frac{e^{-P_o(f)/S_n(f)}}{S_n(f)} \ . \tag{6.11}$$

We assume that there is independent noise in each of $f_{\rm max} t_{\rm obs}$ frequency bins for a given demodulated power spectrum. In general the noise spectra obtained from neighbouring parameter space points will not be statistically independent; however, one may expect that the correlations will be small when the mismatch between the points approaches unity. Therefore we approximate the number of statistically independent noise spectra in our search to be $N_p(t_{\rm obs}, \mu_{\rm max} = 0.3)$. In order that a detection have overall statistical significance α, we must set our detection threshold so there is less than $1 - \alpha$ probability of *any* noise event exceeding that threshold. For a detection to occur the power in the demodulated detector output must satisfy

$$\frac{P_o(f)}{S_n(f)} > \frac{\rho_c}{S_n(f)} = \ln\left[\frac{f_{\rm max} t_{\rm obs} N_p(t_{\rm obs}, \mu_{\rm max} = 0.3)}{1-\alpha}\right] \ , \tag{6.12}$$

where $P_o(f)$ was defined in Eq. (5.1), and ρ_c is the threshold power.

In other words, if the power at a given frequency exceeds ρ_c we can infer that a signal is present; the expected power in the signal is then $\rho_c - S_n$. Thus, the minimum characteristic amplitude we can expect to detect is

$$h_{\rm th} = \sqrt{\frac{(\rho_c/S_n - 1)S_n(f)}{< F_+^2(\Theta, \Phi, \Psi) > (1 - < \mu >) t_{\rm obs}}} \ , \tag{6.13}$$

where $< F_+^2(\Theta, \Phi, \Psi) >$ is the square of the detector response averaged over all possible source positions and wave polarizations. $< \mu >$ is the expected mismatch for a source whose signal parameters $\vec{\lambda}$ lie within a given patch, assuming that all parameter values in that patch are equally likely. We note that the characteristic detector sensitivities $h_{3/\rm yr}$ in Fig. 1 are obtained from this expression by setting $t_{\rm obs} = 10^7$ seconds, $< \mu >= 0$, and $f_{\rm max} t_{\rm obs} N_p = 1$ in the expression for ρ_c; this agrees with Eq. (1.4).

The optimal search strategy is to choose those values of $t_{\rm obs}$ and $\mu_{\rm max}$ which, for some specified computational power $P_{\rm max}$ and detection confidence α, maximize our *sensitivity* Θ which is defined by

$$\Theta(t_{\rm obs}, \mu_{\rm max}) \equiv \frac{1}{h_{\rm th}} \propto \sqrt{\frac{[1 - \frac{s+2}{s+4}\mu_{\rm max}]T}{\rho_c/S_n - 1}} \ , \tag{6.14}$$

where ρ_c/S_n is given by Eq. (6.2). Assuming an overall statistical significance of $\alpha = 0.99$, we have computed the optimal observation time $t_{\rm obs}$ and optimal maximum mismatch $\mu_{\rm max}$, as functions of computing power, for the two searches considered in the previous subsection. The results are shown in Fig. 8.

VII. COMPUTATIONAL REQUIREMENTS FOR A DIRECTED SEARCH

In sections V and VI we examined the computational requirements of an all-sky pulsar search. In this section we examine the computational requirements for a directed pulsar search, by which we mean a search where the position is known but the pulsar frequency and spindown parameters are unknown. Obvious targets in this category are SN1987A, nearby supernova remnants that do not contain known radio pulsars, and the center of our galaxy. Such searches will clearly be among the first performed once the new generation of gravitational wave detectors begin to come on line.

Our treatment of directed pulsar searches closely parallels that of of the all-sky search, so we can be brief. Since the source position (n_x, n_y) is known, we can simply remove the Earth's motion from the data. Below we imagine that the signal has already been transformed to the solar system barycenter. Then the unknown parameters describing the pulsar waveform are

$$(\lambda^0, \lambda^1, \ldots \lambda^s) = (f_0, f_1, \ldots, f_s) \ . \tag{7.1}$$

where the f_i are the same as defined in Eq. (2.1) and s is just the number of spindown parameters included in the frequency model. We again calculate the metrics g_{ij} and γ_{ij} using Eqs. (5.13) and (5.16) respectively, and then calculate N_p using (5.17) (except the integral is now over s-dimensional parameter space). Assuming hexagonal packing in two dimensions and cubic packing in the others, the size of each patch is $V_{\rm patch} = (3\sqrt{3}/4)(4\mu_{\rm max}/s)^{s/2}$. (Except for $s = 1$, where $V_{\rm patch} = 2\mu_{\rm max}^{1/2}$.) We arrive at the expression

$$N_p \simeq \max_{s \in \{1,2\ldots\}} [\mathcal{N}_s G_s(t)] \ , \tag{7.2}$$

where

$$\mathcal{N}_s = \left(\frac{f_0}{1\text{kHz}}\right)^s \left(\frac{40\text{Yrs}}{\tau}\right)^{s(s+1)/2} \left(\frac{0.3}{\mu_{\max}}\right)^{s/2} \tag{7.3}$$

$$G_1(t_{\text{obs}}) = 1.5 \times 10^3\, T^2 \tag{7.4}$$

$$G_2(t_{\text{obs}}) = 6.97 \times 10^1\, T^5 \tag{7.5}$$

$$G_3(t_{\text{obs}}) = 2.89 \times 10^{-4}\, T^9, \tag{7.6}$$

where T is the observation time measured in days. Comparing these results with Eqs. (6.2)–(6.6), we see that for our fiducial parameter values ($f_0 = 1$kHz, $\tau_{\min} = 40$Yrs, $\mu_{\max} = 0.3$) and observation times T of order a week, N_p is $\sim 10^5$ times larger for an all-sky search than for a directed search. Another way of putting this is: after using one's freedom to adjust the frequency and spin-down parameters in optimizing the fit, only $\sim 10^5$ distinguishable patches on the sky remain. Equivalently, a single directed search can cover an area of $\sim 10^{-4}$ steradians. Thus ~ 1000 week-long, directed searches would be sufficient to cover the galactic center region.

We can calculate the optimal $\mu_{\max}$ and t_{obs} as a function of computing power for a directed search in the same way as we did for the all-sky directed search. (Except the factor $\frac{s+2}{s+4}$ in Eq. 6.14 becomes $\frac{s}{s+2}$ for the directed-search case.) The results are shown in Fig. 9, for our two fiducial types of pulsar. We see that knowing the source position in advance increases t_{obs} by only a factor of ~ 10, for 1 Tflops computing power. The resulting gains in sensitivity can be seen in Fig. 2.

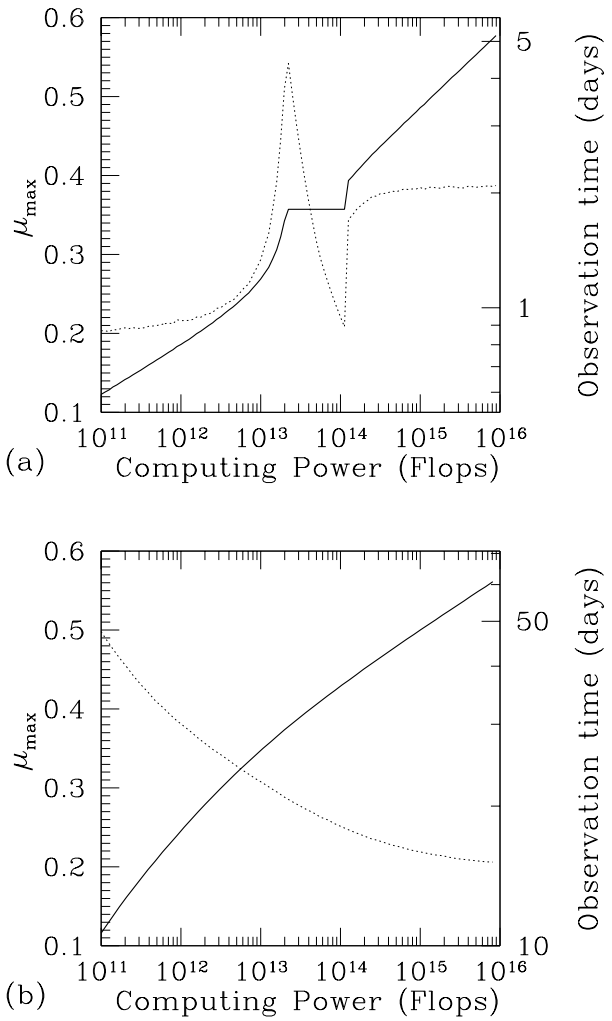

FIG. 8. The optimum observation time (thick solid line), and maximal projected mismatch (dotted line) as functions of available computational power. Both graphs assume a threshold which gives an overall statistical significance of 99% to any detection (although the results should be insensitive to the precise value). Each of the graphs corresponds to: (a) the situation encountered when searching for periodic sources having gravitational wave frequencies up to 1000Hz, with minimum spindown ages $\tau_{\min} = 40$Yrs. (b) The equivalent results for gravitational wave frequencies up to 200Hz, with minimum spindown ages $\tau_{\min} = 10^3$Yrs. The transition region seen in figure (a) is due to the fact that a longer integration time would require searching over an additional spindown parameter, as seen in Fig. 6. In this region it is more efficient, as one adds computational power, to lower mismatch thresholds, rather than searching over the additional parameter.

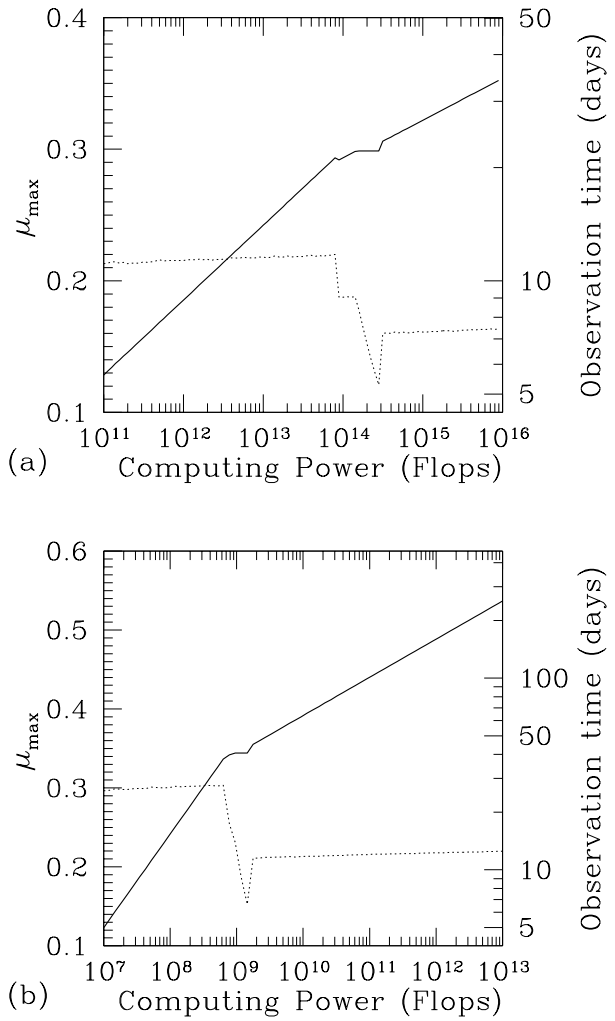

FIG. 9. The optimum observation time (thick solid line), and maximal projected mismatch (dotted line) as functions of available computational power for directed searches. Both graphs assume a threshold which gives an overall statistical significance of 99% to any detection (although the results are insensitive to the precise value). Each of the graphs corresponds to: (a) the situation encountered when searching for periodic sources having gravitational wave frequencies up to 1000Hz, with minimum spindown ages $\tau_{\rm min} = 40$Yrs. (b) The equivalent results for gravitational wave frequencies up to 200Hz, with minimum spindown ages $\tau_{\rm min} = 10^3$Yrs. The transition regions, where the optimum observation time does not increase, are due to the fact that a longer integration time would require searching over an additional spindown parameter.

VIII. FUTURE DIRECTIONS

Searching for unknown sources of continuous gravitational waves using LIGO, or other interferometers, will be an immense computational task. In this paper we have presented our current understanding of the problem. By applying techniques from differential geometry we have estimated the number of independent points in the parameter space which must be considered in all-sky and directed searches for sources which spin down on timescales short enough to produce observable effects; these numbers were used to compute the maximum achievable sensitivity for a coherent search (see Fig. 2). Furthermore, the metric formulation can be used to optimally place the parameter space points which must be sampled in a search.

Our analysis takes no account of bottlenecks in the analysis process due to data input/output and interprocessor communication. These are important issues which may impose further constraints on the maximum observation time; however, it seems premature to address such problems until we know the hardware that will be used to conduct searches for continuous waves.

Unfortunately, Fig. 8 shows that it will be impossible to search, in one step, 10^7 seconds worth of data over all sky positions. However it is also unnecessary. We foresee implementing a hierarchical search strategy, in which a long data stream is searched in two (or more) stages, trading off sensitivity in the first stage for reduced computational requirements. Having determined a number of potential signals in the first stage—presumably at a threshold level which allows many false alarms due to random noise—these candidate events would be followed up in the second stage, using longer integration times. The longer integration times would be possible because the search would only have to be performed over much smaller regions of the parameter space, in the neighbourhoods of the candidate signal parameters. In this way, one can achieve a greater sensitivity than a coherent search using the same computational resources.

Clearly one can imagine many different implementations of this rough strategy, and we have not yet determined the optimal one. Nevertheless, we have considered the simple example where the data is searched in two stages. Candidate signals from an all-sky search of a short stretch of data [$T^{(1)}$ seconds long] are followed up using longer Fourier transforms to achieve greater sensitivity. One can estimate $T^{(1)}$ using Fig. 8 and an assumption that roughly half of the total computing budget is used on the first stage; this turns out to be a valid assumption. A simple argument along these lines goes as follows. Consider a search for 'young, fast' pulsars that begins by coherently analyzing stretches of data that are all ~ 1 day long (possible with $\sim 4\times10^{12}$ flops, by Fig. 8). Imagine that in the second stage of the search one follows up all templates such that $P_0(f,\vec{\lambda}) > 4.6 S_n(f)$, by seeing whether templates with roughly the same parameter values are exceeding this threshhold every day. (Here $P_0(f,\vec{\lambda})$ is the power of the stretched data at frequency f, for stretch $\vec{\lambda}$. This threshhold implies that one is following up only one out of every hundred templates.) It seems likely that this second stage will not be more computationally intensive than the first. To exceed this thresh-

old, a pulsar must have $h_c \gtrsim 12 h_{3/\mathrm{yr}}$. This is factor of roughly 3 better than if one restricted oneself to coherent searches considered above, but is a factor of 3 worse than the sensitivity one could achieve with unlimited computing power.

A refinement of this strategy would be one in which the first pass consists of several incoherently-added power spectra. That is, one slices the data into N sequential subsets, performs a full search (as described in this paper) for each subset, and adds up the power spectra of the resulting searches for each of the parameter sets. This technique has been used to good effect by radio astronomers searching for pulsars [28]. Since the addition of power spectra is incoherent, there is a loss of signal-to-noise ratio in the final summed power spectrum of $1/\sqrt{N}$ in relation to a full coherent search over the whole timescale. However, the computational savings involved allow one to search stretches of data which are much longer overall. For some optimal choice of N, this will result in higher sensitivities when one follows up candidate detections using coherent searches. D. Nicholson (private communication) has estimated that a 1Tflops computer could perform such a search of 10^7s of data, over all sky positions but ignoring pulsar spindowns. A subsequent paper will present a concrete analysis of this and other hierarchical scenarios [34].

ACKNOWLEDGMENTS

PRB and TC wish to thank Kip Thorne for suggesting that we work on this problem; we are extremely grateful to him for his encouragement and suggestions. PRB and TC also thank Bruce Allen, Kent Blackburn, Jolien Creighton, Sam Finn, Scott Hughes, Ben Owen, Massimo Tinto, and Alan Wiseman for useful discussions. This work was supported in part by NSF grant PHY-9424337. PRB is supported by a PMA prize fellowship at Caltech. BFS wishes to thank Gareth Jones, Andrez Krolak, Andrew Lyne and David Nicolson for useful conversations. CC thanks Lee Lindblom for helpful discussions; CC's work was supported in part by NSF grant PHY-9507686 and by a grant from the Alfred P. Sloan Foundation.

[1] A. Abramovici *et al.*, Science **256**, 325 (1992).
[2] C. Bradaschia *et al.*, Nucl. Inst. A. **289**, 518 (1990).
[3] L. Blanchet, B. R. Iyer, C. M. Will, and A. G. Wiseman, Class. Quant. **13**, 575 (1996).
[4] T. A. Apostolatos, Phys. Rev. D. **54**, 2421 (1996).
[5] B. Owen, Phys. Rev. D. **53**, 6749 (1996).
[6] R. Balasubramanian, B. S. Sathyaprakash, and S. V. Dhurandhar, Phys. Rev. D. **53**, 3033 (1996).
[7] K. S. Thorne, in *Three hundred years of gravitation*, edited by S. W. Hawking and W. Israel (Cambridge University Press, Cambridge, 1987), Chap. 9, pp. 330–458.
[8] E. E. Flanagan, Phys. Rev. D. **48**, 2389 (1993).
[9] K. Compton, Ph.D. thesis, University of Wales, Cardiff, 1996.
[10] D. Nicholson *et al.*, Physics Letters **A 218**, 175 (1996).
[11] B. Allen, in *Proceedings of the Les Houches School on Astrophysical Sources of Gravitational Waves*, edited by J. A. Marck and J. P. Lasota (Cambridge University Press, Cambridge, 1996).
[12] S. Bonazzola and E. Gourgoulhon, Astron. Astr. **312**, 675 (1996).
[13] S. Chandrasekhar, Phys. Rev. Let. **24**, 611 (1970).
[14] J. L. Friedman and B. F. Schutz, Ap. J. **222**, 281 (1978).
[15] R. V. Wagoner, Astrophys. J. **278**, 345 (1984).
[16] M. Zimmermann and E. Szedenits, Jr., Phys. Rev. D. **20**, 351 (1979).
[17] R. D. Blandford, unpublished (1984).
[18] B. F. Schutz, in *The Detection of Gravitational Waves*, edited by D. G. Blair (Cambridge University Press, Cambridge, 1991), Chap. 16, pp. 406–451.
[19] J. C. Livas, Ph.D. thesis, Massachusetts Institute of Technology, 1987.
[20] G. S. Jones, Ph.D. thesis, University of Whales, 1995.
[21] T. M. Niebauer *et al.*, Phys. Rev. D **47**, (1993).
[22] S. L. Shapiro and S. A. Teukolsky, *Black Holes, White Dwarfs and Neutron Stars* (Wiley, New York, 1983), Chap. 10.
[23] A. G. Lyne, Phil. Trans. R. Soc. Lond. A **341**, 29 (1992).
[24] R. Narayan and J. P. Ostriker, Astrophys. J. **352**, 222 (1990).
[25] R. N. Manchester, Phil. Trans. R. Soc. Lond. A **341**, 3 (1992).
[26] S. R. Kulkarni, Phil. Trans. R. Soc. Lond. A **341**, 77 (1992).
[27] A. Wolszczan *et al.*, Nature **337**, 531 (1989).
[28] S. B. Anderson, Ph.D. thesis, California Institute of Technology, Pasadena, California, 1993.
[29] M. Zimmermann, Nature **271**, 524 (1978).
[30] M. Zimmermann, Phys. Rev. D. **21**, 891 (1980).
[31] D. V. Gal'tsov, V. P. Tsvetkov, and A. N. Tsirulev, Sov. Phys. —JETP **59**, 472 (1984).
[32] J. Middleditch, Ph.D. thesis, University of California, Berkeley, 1975.
[33] M. Tinto, private communication (1997).
[34] P. R. Brady and T. Creighton, paper in preparation (1996).